

A Methodology for Extracting Objective Color From Images

Mark W. Powell, Sudeep Sarkar, *Member, IEEE*, Dmitry B. Goldgof, *Senior Member, IEEE*, and Krassimir Ivanov, *Member, IEEE*

Abstract—We present a methodology for correcting color images taken in practical indoor environments, such as laboratories, factories, and studios, that explicitly models illuminant location, surface reflectance and geometry, and camera responsivity. We explicitly model surfaces by taking our color images with corresponding registered three-dimensional (3-D) range images, which provide surface orientation and location information for every point in the scene. We automatically detect regions where color correction should not be applied, such as specularities, coarse texture regions, and jump edges. This correction results in objective color measures of the imaged surfaces. This kind of integrated, comprehensive system of color correction has not existed until now. i.e., it is the first of its kind in computer vision. We demonstrate results of applying this methodology to real images for applications in photorealistic rerendering, skin lesion detection, burn scar color measurement, and general color image enhancement. We also have tested the method under different lighting configurations and with three different range scanners.

Index Terms—Burn scars, color calibration, color correction, computer vision, image processing, light calibration, range images, realistic re-rendering, skin color.

I. INTRODUCTION

IT IS VERY OFTEN desirable to use information in color images to solve problems. However, the color image formation process is very complicated, even in constrained environments such as an indoor laboratory, factory, or studio where illumination can be engineered and color cameras can be modeled. In medical imaging applications such as burn scar characterization and skin lesion detection, issues of surface shading, illumination variation, and surface texture variegation complicate most automated analysis algorithms. In the context of imaging for computer graphics, i.e., capturing texture map images, surface shading, texture and illumination variation effects have to be modeled by the rendering process. We would like to elimi-

nate these effects in images that we capture for use as texture maps so that our rendering will have more realism. Our motivation in this research is to use light source calibration and surface geometry estimation to extract *objective color from images* that have varying illumination [1], [2]. For our system we assume that the scene is an indoor scene that is characterized as having one or more illuminant that may be practically considered as point sources, or groups of point sources. The surfaces in the scene may be of any shape and may have surface reflectance that can be modeled by a bidirectional reflectance distribution model (BRDF) at any given surface point or patch, i.e., having a diffuse and specular component. We assume that for application a conventional color charge coupled device (CCD) camera is used. The proposed system is suitable for any application where camera position, response, and illumination location may be modeled prior to imaging one or more scenes in the local environment.

Our strategy for color correction of images is based on modeling the three major components of image formation: illumination, reflectance, and responsivity. The integral of the product of illumination, reflectance, and responsivity with respect to wavelength yields color stimulus. We model illumination by calibrating the positions of all of the illumination sources that light the scene. We use range images to measure the geometry of the scene and correct for surface shading and brightness variation due to changing distance from the illumination. We also use range information to identify specular highlights, edges, and regions of coarse surface texture. Finally, we model the position, orientation and responsivity of the camera used for image capture.

Fig. 1 illustrates the stages of the objective color extraction methodology. Each aspect of the scene (illumination, surface information, and camera response) are modeled and calibrated. Before imaging the scenes of interest for a given application, we calibrate light sources using two devices: a Macbeth Color checker chart, and a triangulation device. The color chart (not shown in this print version) characterizes camera response and illumination color. The triangulation device gives us the three-dimensional (3-D) locations of the light sources relative to the camera. The triangulation device used for calibrating the light source direction and location consists of three specular spheres of known size and relative location. Our technique for calibrating the light source geometry is highly practical for a laboratory or studio setting since it assumes that the light sources are near the scene of interest. We estimate scene geometry using range imagery to estimate surface normals. We then segment the range image and use range texture to identify smooth or coarse surfaces in the scene. Texture from range, especially from general surfaces is a new innovation in information extraction from

Manuscript received November 4, 2002; revised July 14, 2003 and December 16, 2003. This work was supported in part by the Whitaker Foundation Biomedical Engineering Research Grant and in part by the National Science Foundation under Grants EIA-9729904 and EIA-0130768. This paper was recommended by Associate Editor P. Bhattacharya.

This paper has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors. This includes .jpg files of all color images and an MPEG animation to go along with Fig. 20. This material is 856 KB in size.

M. W. Powell is with the Mobility Systems Concept Development Section, Jet Propulsion Laboratory, Pasadena, CA 91109 USA (e-mail: mpowell@csee.usf.edu).

S. Sarkar and D. B. Goldgof are with the Department of Computer Science and Engineering, University of South Florida, Tampa, FL 33620 USA (e-mail: sarkar@csee.usf.edu; goldgof@csee.usf.edu).

K. Ivanov is with Verizon Services Group, Tampa, FL 33637 USA (e-mail: kivanov@csee.usf.edu).

Digital Object Identifier 10.1109/TSMCB.2004.832177

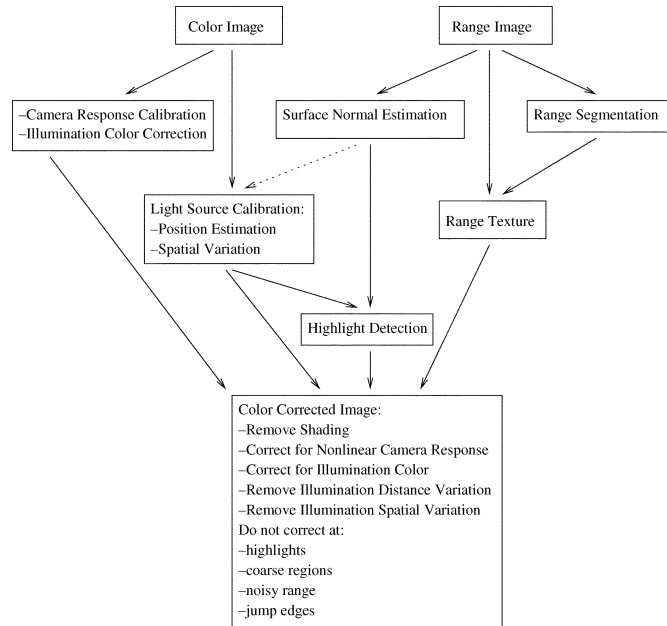


Fig. 1. System overview diagram for computing objective color.

range images developed for this work. Next, we combine the light source locations and scene geometry information to estimate where specular highlights appear in the image. Specular highlight regions contain mainly the intensity of the illuminant, and not the underlying surface, thus, no objective surface color information is present in these areas. We also detect coarse texture regions, noisy range, and jump edges. These types of surface regions violate our aforementioned surface modeling assumptions. Lastly, we correct for nonlinear camera response and off-white illumination color. This kind of integrated, comprehensive system of color correction has not existed until now, i.e., it is the first of its kind in computer vision.

The following subsection gives a discussion of related work. Section II describes the camera calibration, light source calibration, color correction methodologies. The results of applying these techniques to several applications including burn scars, photorealistic rerendering, and color correction will be presented in Section III. Finally, Section IV offers conclusions drawn from this research.

A. Related Work

One class of related work is the formulation of illumination-invariant measures of objective color [3]–[5]. We choose instead to directly calibrate the illumination and surface geometry of an imaged scene. This approach mirrors the computer graphics paradigm of modeling the scene and illumination to produce a realistic image, except that we seek to extract accurate surface reflectance parameterization from actual scenes.

Another area of related work is that of computational color constancy. Strictly speaking, humans do not have color constancy in the sense that the exact color of an object is recognizable under any arbitrary illumination conditions, as a variety of psychophysical experimental evidence demonstrates [6], [7]. The area of computational color constancy seeks to achieve through algorithmic means what human color recognition cannot: color recognition under arbitrary lighting conditions. This is a difficult and under-constrained problem,

as evidenced by the large number of attempts to solve it, for example [7]–[16]. Such solutions often make unrealistic assumptions about the imaged scene that are impractical, such as two-dimensionality and smoothly-varying illumination. Our approach mitigates the under-constrained characteristics of the problem by taking advantage of range data and calibrated illumination.

One particularly relevant related project is the Digital Michelangelo Project [17]. This project created a sophisticated system for recording the surface reflectance and geometric shape of statues created by Michelangelo. A series of color and range images were obtained over the entire outer surface of a number of marble statues which were converted to surface reflectance information using a calibrated camera, illumination, and range information from the surface. In order to acquire reliable color, the spotlight illumination was carefully chosen and controlled with respect to the camera, and ambient illumination effects were later subtracted out in post processing. Like the color processing pipeline used in this project, our system estimates the diffuse component of the bi-directional reflectance distribution function (BRDF) [18] over the entire surface by taking advantage of the coregistered surface geometry. Our methodology for color correction differs from this approach by calibrating all of the contributing light sources *a priori* and using a single image from which surface reflectance characteristics are derived.

In the area of texture from range, there are other approaches related to ours that are of interest. One approach described in [19] uses co-occurrence matrices as a means of deriving texture measures from range. They also resample the range data and use a robust least median squares estimator to remove outliers. However, this approach has only been demonstrated on images from one structured light scanner where the sampled surfaces were orthogonal to the optical axis of the camera. In our results we demonstrate applicability of our approach to both planar and nonplanar surfaces that may vary in surface orientation with respect to the direction of view. In [20] there is another co-occurrence matrix based approach for measuring 3-D texture. The quantitative results were given only for images of simple scenes with planar surfaces, and the authors note that when the method is applied to real scenes with surfaces of varying type and orientation, it is only practical to produce qualitative results. In our results we have included considerable quantitative results from our approach to texture from range. Another interesting related work is [21], in which a simple histogram-based algorithm is applied to laser range finder images taken from a mobile robotics platform to quickly (in real time) classify regions of foliage. Although this approach was effective for classifying regions of grass and rocks, it lacks the fidelity needed for the applications that are demonstrated in our results, i.e., photorealistic rerendering and skin lesion detection.

II. METHODOLOGY

A. Camera Response Calibration and Illumination Color Correction

When color images are captured by a digital camera, such as a CCD, the response of the three color bands are usually somewhat different from each other. For instance, one band may respond more strongly to intensity than another. If this is the




Fig. 2. Plot of relative intensity versus camera response. The curves are based on five samples for each color band. The white line represents the ideal camera response. Actual curves shown as green (highest), red, and blue (lowest).

case, then the resulting image will have a hue shift to that color when imaging a white or grey patch, e.g., if the blue channel is stronger than the red and green, the white patch will appear to have a bluish hue. Also, the responses of a CCD are generally nonlinear with respect to illumination. Many machine vision algorithms implicitly assume that camera response varies linearly with intensity, such as in edge detection that is based fundamentally on contrast between groups of pixels. We would like to calibrate the camera response in order to measure color from images objectively.

Our camera response calibration uses the Macbeth Colorchecker reference chart. This chart contains 24 patches that resemble colors from the natural environment (skin, grass, soil), primary colors (red, green, and blue), secondary colors (yellow, cyan, and magenta) and achromatic colors ranging from white to black. We use the achromatic, neutral color patches to sample camera response. For all six neutral patches, we take the mean intensities in each band over a 10×10 window centered on the patch. We know the absolute reflectance values of the patches. These patches also have the useful property of reflecting light the same way in all parts of the visible spectrum. The six neutral patches have reflectance values (expressed as percentage of illumination intensity) of 2.5% (black), 9.0%, 21.0%, 36.5%, 59.5%, and 88.5% (white).

Fig. 2 shows a plot of the reflectance values from an image of the chart with estimated curves drawn through the points. Five of the six patches were actually used to create this plot (black was omitted). The goal of this calibration is to remap the actual responses (estimated by the curves in this figure) to the ideal linear response (represented by the line in the figure). There are three curves here, one for each color channel of the camera (not shown in this print version) used for our examples: green (highest), red, and blue (lowest). Since we only start with five samples per band, we would like to interpolate a continuous function to represent the camera response over the entire range of the camera. In this example we used a general order polynomial interpolation algorithm, and we found that a fourth order polynomial fit produces the smoothest, best fitting curve approximation to our camera's measured response. For other cameras in general, we recommend starting with a second order fit and

increasing order as needed to maximize curve smoothness and minimize residual fit error to measured response values. We found that including the black patch response sometimes produced curve fits that were not smooth. To improve the fit, we removed the response from the black patch from the fitting algorithm inputs. The black patch is the least important measurement, since there the sampling of responses is relatively dense in the lower intensities compared to the higher intensities. When the polynomial approximation of the camera response has been estimated, a lookup table can be formed that remaps actual response to ideal response over the whole range of intensities from $0 \dots 255$. When this lookup table transform is applied to a color image, the response is linearized. Another important result of this procedure is that the camera responses for all three color are all rescaled so that white patches are properly white-balanced. In other words, this calibration procedure compensates for the camera response and the color of the illumination at the same time. As we can see in Fig. 3, not only is the overall contrast of the image enhanced from the linearization, but the hues in the image have been corrected for the yellowish tungsten illumination used to create this image.

Other concerns to keep in mind when acquiring color images are the response range and clipping effects. A given camera will have a given response range for a particular aperture setting on its lens and the particular gain (response amplitude) setting. These settings should be carefully optimized for the imaging environment. If the response range is too small (not enough gain or not enough light on the sensor plane through the lens aperture), then the signal to noise ratio of the image information will be quite low. However, because the upper and lower bounds on the camera responses are finite, setting the gain too high or opening the aperture too much may result in clipping of very bright regions. Clipping means that very bright or very dark regions have no useful information because their intensities are outside the response range of the sensor. An optimal response setting can be achieved by viewing the histogram of the color image of the Colorchecker chart and ensuring that it occupies at least two-thirds of the overall response range of the camera while making sure that the histogram does not contain any high values near either extreme end of the range (0 or 255 for an 8-bit intensity image). Finding the optimal range of response for our work involved an iterative process of varying the aperture setting of the camera lens and the gain setting of the camera sensor, examining the histogram of an image of the reference chart at each point, and evaluating by inspection what settings maximized the range of intensity response while avoiding clipping of bright surfaces. For color cameras in general, it may be desirable to independently tune the gain and aperture independently for each color channel. However, since using different gain for each color response independently would also rescale the response values, it would be necessary to correct for these variations to white-balance the image. For our work, we were not able to practically achieve independent gain and aperture tuning for each color channel since this process involved the physical manipulation of the camera hardware. Physically adjusting the camera created problems for maintaining the position and orientation of the camera relative to the structured light scanner used to acquire range information that is used in later processing stages of our method. We



(a)



(b)

Fig. 3. (a) Image of a person's hands under tungsten light before color response calibration. (b) After calibration the image response is linearized, increasing the contrast and correcting for the illumination color.

therefore converted our color images to grey scale and tuned the camera settings for the combination of all three channels.

Since the brightest patch on the Colorchecker chart only has about 88.5% reflectance, using only this chart might still result in clipping when sampling shiny, bright objects like pale, oily skin. To properly set the response range for scenes where there are bright, shiny surfaces, an object of the expected type should be placed in the scene with the chart to augment the histogram-based inspection.

B. Spatial Distribution of the Illumination Field

Very often a significant source of intensity variation in an image is from the variation of the illumination strength as a function of the position in the image plane. We call this the variation of the spatial distribution of the illumination field. We assume that light sources we will encounter are point sources, or other configurations that can be approximated by one or more point light sources.

For our discussion, we will first consider a single illuminant: a GretagMachbeth Sol-Source daylight lamp. This lamp is essentially a small halogen bulb within a reflective hemispherical housing with an aperture that produces a roughly conical volume of illumination. The field has a finite boundary, and we wish to characterize the illumination strength as a function of position within the illumination volume. We chose an experimental observation approach, where we placed a white planar surface at varying distances from the light source and observed the resulting intensity strength. Fig. 4 shows two examples of the illumination field on the white plane oriented perpendicularly to the direction of illumination with 20-cm difference in distance between the two images.

If we consider the horizontal line that passes through the center of the brightest spot in the center of the illumination field over a set of images, we obtain a profile of the illumination intensity as a function of distance from the source and the position in the image plane. Experimentation with the resultant intensity curves shows that they may be well-approximated by a Gaussian




Fig. 4. Two examples of the illumination field on the white plane oriented perpendicularly to the direction of illumination with 20-cm difference in distance between the two images. Note: the outer case of the lamp is visible at the top of the image.

(see Fig. 5). With this assumption, the intensity at a location (i, j) in an image can be modeled by the following equation:

$$I_{\text{surface}}(i, j, \lambda) = A(d)I_0(\lambda)e^{-\theta^2/\sigma_l^2} \quad (1)$$

where I_0 is the intensity of the light source at a wavelength of λ , A accounts for camera gain and is also a function of d , the distance from the surface point that is imaged at pixel (i, j) and the light source, θ is the angle from the axis of the illumination cone, and σ_l captures the extent of angular spread of the light.

We performed a function minimization procedure to calculate the best estimates of the parameters of the distributions at various distances. We found that the term σ_l in (1) above increases linearly with distance from the light source according to

$$\sigma_l = 0.1775d + 44.575 \quad (2)$$

where d is the distance from the light source to the planar surface measured in millimeters and d is less than one meter. At a distance of a meter or more, the value of sigma remains constant. We also estimated the function $A(d)$, which increases nonlinearly with distance according to

$$A(d) = \frac{11\,400\,000}{(d - 170)} \quad (3)$$

where d is the distance from source to surface as before, in millimeters. We also found that for the purposes of experimentation, the spatial variation of illumination strength is negligible




Fig. 5. Plot of reflected intensity strength over the image plane at various distances from the daylight lamp.

at distances of 1 m or greater for this light source. This simplification allows us to eliminate this step in color correction under these conditions. Otherwise, calibration of the illuminant position and orientation would be necessary to apply a linearization to compensate for this effect in the image, as we do in our next example.

As a multiple illumination scenario, we consider the commercially-available Cyberware 3030 HRC scanner which has four halogen bulbs arranged in a vertical line. These positions remain fixed with respect to the camera, but the object moves in the direction of the horizontal axis of the image plane, and each image sample is taken in only a single vertical column as the object moves. Therefore, there is no spatial variation in illumination in the image in the horizontal direction at all. The remaining spatial variation in the vertical direction can be expressed at a one-dimensional function of height in the scene. To correct this variation, we use basically the same idea as with the single daylight lamp. We measured the intensity change in the vertical direction again using a white plane at different depths from the camera. The resulting intensity variation curve is recorded and converted from a general curve (Fig. 5) to deviations from a straight line that is horizontal and passes through the highest point on the curve. The difference from the straight line is converted to a relative value ranging from 0 to 1. To correct the image, each pixel is brightened according to the relative brightness from its associated height. As a result, every point in the image becomes as bright as the point with the brightest illumination.

C. Light Source Calibration

We can use the locations of the light sources that create the image to model effects such as shading and illumination fall-off due to increasing distance. In order to calibrate these locations




Fig. 6. Light source calibration setup illuminated by three different light sources simultaneously. The three spheres are specular on one side (shown on the front-most two spheres) and Lambertian on the other (as shown on the rear-most sphere). The 3-D positions of the specularities visible on the surface are used to triangulate for the light source positions.

explicitly, we have devised a novel calibration device. The device consists of three spheres (see Fig. 6). The locations of the specularities that are imaged by the camera are produced by the illuminant we wish to locate. Each pair of specularities that correspond to a single light source will allow us to triangulate its position. This work is summarized here, but it described in greater detail in [2], [22].

For many of our results we acquire images using a Cyberware 3030 HRC scanner to image registered color and range. In the case of the Cyberware scanner, the four lights that illuminate the scene are at fixed positions relative to the camera. However, the camera itself moves horizontally in relation to the scene, so

illumination positions are only constant for each vertical column of pixels in the image. When using knowledge of illumination position to correct color, each column of the color image must be processed with respect to the illumination position at the time it was captured. The light calibration methodology can still be applied here, however.

D. Color Correction

We can use light source calibration to remove the shading on surfaces in the image which get darker in intensity as the angle between the surface normal and the optical axis increases. To arrive at intrinsic color of surfaces, we pre-process the color image and adjust the intensity values based on the light source direction. The amount of light reflected back from the surface to the camera depends on the reflectance properties of the surface, or albedo. We assume a Lambertian surface, where the reflected light varies directly with the cosine of the angle, θ , between the direction of incoming light and the normal of the surface at a given point. We estimate the surface normal at each pixel using a range image and a plane-fitting method based on the covariance matrix of a local region. This algorithm is given in detail in [1]. For every point in the image, the intensity value is divided by $\cos(\theta)$. Dividing the intensity values by this directionality factor results in an image that is very close to how the surface would appear if the entire surface were flat and oriented perpendicularly with respect to the direction of illumination.

E. Detecting Where Correction Should not be Applied

An underlying assumption in our color correction strategy is that the surface normal is changing smoothly everywhere in the image. There are many instances where this is not the case, such as regions representing edges, noise, and coarse surfaces. Any region in the image where the surface normal cannot be estimated well will result in an over-correction or under-correction in the color.

Edges are a problem, since they usually represent a situation where surface orientation changes dynamically. Jump edges are depth discontinuities where two different surfaces meet. If the pixel overlaps two different surfaces, it will generally take on a depth value somewhere in between the depths of the two surfaces. This usually makes these pixels detectable via simple thresholding, comparing the edge pixels to neighboring pixels. Crease edges are a more subtle case. Here, two different surfaces are joined at an edge where the depth does not change in any significant way, and it is only the surface type or orientation that changes. We may be able to detect this type of edge by searching for where the surface normal changes suddenly at the edges between two surfaces. If we perform edge detection before we correct the color, we can create a mask of pixel positions where we do not want to correct color. The information in the mask can also be used at higher levels of vision processing.

Noisy range pixels are a problem also, since, like edges, they make surface normal estimation unreliable. Certain types of noise can be reported by the scanner as void pixels. In other cases, a noisy range reading will appear as an outlier compared to the majority of the nearby legitimate data. These types of outliers may be detected by a manual thresholding procedure.

If we add this information to the mask along with the edges, we can avoid over-correcting the color in the resultant image where noise exists. Other types of noise, which may lie close to the surface, can not be detected in a straightforward way, but may be detected by computing range texture as we discussed in Section II-E2. Range texture computation will allow us to detect coarse regions where surface normal estimation does not vary smoothly over the region.

1) *Specular Highlights*: Another assumption that we have made until now is that the color of the pixel in the original image is due to its body reflectance. This is only true if the surface reflectance is matte. If the surface is shiny, then certain regions on the surface will contribute specular reflectance that is seen by the camera. This is especially the case where there are specular highlights. The color at these pixels is most like the illuminant itself and not the underlying surface. At these regions, there is not enough body reflectance for us to be able to correct the color. However, if we can determine where highlights are likely to be found, we can mask them out when performing color correction, and report them as not useful for identifying the body reflectance properties of the underlying surface.

We can use light source calibration information and range to test at every pixel whether the surface geometry will support a specular reflection. Using a range image, we can estimate the surface normal and the camera position. We can also calibrate the location of each light source, as shown in Section II-C. The law of reflection requires that the incoming illumination vector and outgoing reflectance vector have equal angles with respect to the surface normal. Additionally, since we are working with 3-D geometry, we must ensure that the illumination vector, reflectance vector, and surface normal are coplanar. If this relationship is satisfied, and the surface is specular, a highlight will occur at the pixel in question. Setting a threshold on the equality of the incoming and outgoing angles will let us detect this condition at each pixel in the image. Also, we can test whether the three vectors involved are coplanar by taking the cross product of two of the vectors and then the dot product of the resultant vector with the remaining vector. If the three vectors are coplanar, then the dot product will be zero. Thresholding the proximity of this dot product to zero will allow us to detect it. When the two thresholds are satisfied simultaneously, the geometry will support a specular highlight. If the surface is indeed shiny, then there will be a highlight at these locations in the image.

Properly selecting the thresholds just described is implementation-specific in several ways. The surface normals must be estimated using the available range data, and the quality of the 3-D surface information may vary with surface type (smoothly varying versus highly undulating) and ranging accuracy. Outliers that are often present in range data may be tolerated or not depending on the quality of the implementation of surface normal estimation. The choice of the thresholds can also be related to the extent of the specular lobe one wants to label. Another less significant but still appreciable factor is the nature of computing with floating point numbers, which prevents direct comparison of the incoming and outgoing angles of reflection for equality (a small fractional number should be chosen to compare to the difference of the angles). For our results, a

least-squares planar fit that rejected outliers based on a fit error measure worked best.

2) *Range Texture*: The idea of range texture is to extract features of the spatial distribution of the height or depth of points in the image. This concept is analogous to the idea of extracting texture from intensity images, or computing features from the spatial distribution of grey levels. Some examples of useful features include coarseness, smoothness, regularity, and directionality. Texture features are computed over a neighborhood of pixels, since they are meaningless for a single pixel. Since range data may change with varying positions of objects and camera geometry, we must take care to standardize range data across a set of images.

The data standardization procedure is as follows. We perform a segmentation of the range data in each neighborhood. The range segmentation algorithm, described in detail in [23], first performs edge detection to the range image to establish a preliminary set of regions. Each region is then characterized by surface fitting. First a planar fit is attempted, and if the error in fit exceeds a threshold value, then a nonplanar fit is made using a polynomial surface of up to degree 4. (We tuned the threshold value on the planar fit error based on our empirical results; see [23] for further details.) In this way each region is classified as planar or nonplanar.

We then compute the residual error at each pixel from the fitted surface used to segment it. If the fitted surface is planar, then the residual error is the distance from the measured surface location to the corresponding fitted plane location. Otherwise, the residual error from the polynomial surface fit is computed. Since there is no closed form solution for the residual distance from a point to a fourth order polynomial surface, we take the error in only in the z axis (the direction of depth measurement).

Next, the residual errors of each point from the fitted surface are rescaled to integer values ranging from 0 to n , where n is a maximum value of depth that is empirically determined based on the images under consideration and the resolution of the range capture device. We then normalize the residual values (dividing them by n^2 , the dimensionality of the co-occurrence matrix). The residual values are then input into a co-occurrence matrix. Once the co-occurrence matrix is formed, a set of second-order statistical feature values, such as energy, entropy, correlation, etc., are computed based on the matrix entries.

F. Algorithmic Complexity

The computational complexity of our approach is quite practical compared to many other image processing algorithms. Let us consider an image of size $m \times n$ in pixels. The camera response correction is implementable as a lookup table, which requires only $O(mn)$ operations. If the light source position(s) are known, the correction for varying distance and spatial distribution of the illumination are also $O(mn)$. Correction for shading and specular detection requires knowledge surface normal estimation, which is a window-based operation, making these algorithms $O(mni)$ where the window size is $i \times j$ pixels. The algorithm that we used for finding edges in range images is scan line based rather than pixel based, making it very efficient. Its complexity is therefore $O(m + n)$. The range image segmenter




Fig. 7. (a) Original image of three yellow blocks. (b) Color corrected image. (c) PCT/MC segmentation of the three-block image. (d) Segmentation of the color corrected image. A supplementary color jpg file is available at <http://ieeexplore.ieee.org>.




Fig. 8. Cyberware images of three southwestern native American style pots. (a), (c), and (e) Original color images, shaded on the right. (b), (d), and (f) Corrected images with shading removed. A supplementary color .jpg file is available at <http://ieeexplore.ieee.org>.

that is used for range texture uses a surface fitting algorithm that is pixel-based, which makes it $O(mn)$.

III. RESULTS

In this section, we show results of the color correction system both individually and as a whole. First, we present the results of

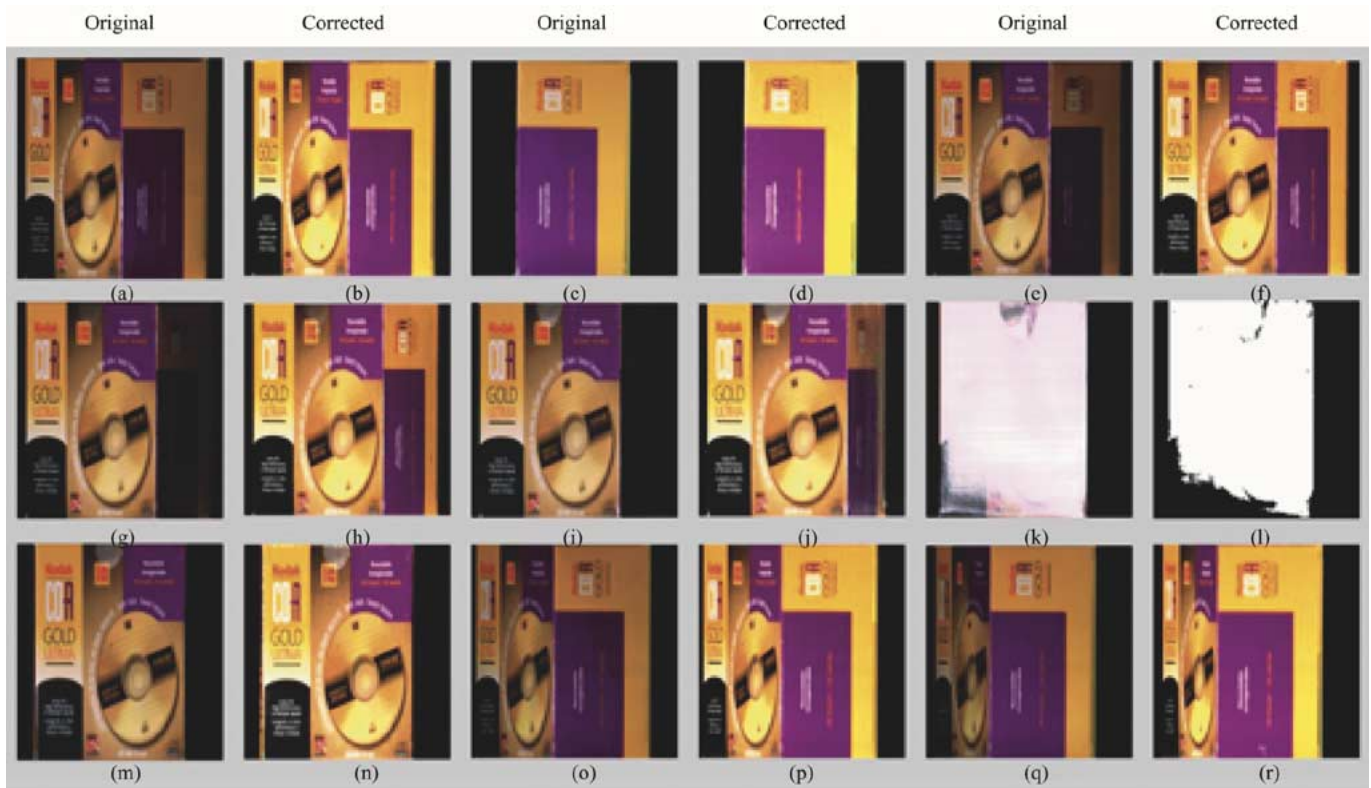


Fig. 9. CD box at various orientations, with correction.

color correction that is applied to images where the various surfaces all conform to our assumptions of smoothness and matte reflectance. We show how color correction can be used in scenes illuminated with a single light source or by as many as four illuminants simultaneously. Next, we present a statistical analysis of the accuracy of the corrected color. We then demonstrate the results of both our specular highlight detection methodology and coarse surface detection methodology which identify regions where color correction should not be applied. All of these individual algorithms are then used together to demonstrate their use as an integrated system for surface analysis and color correction.

Next, we show how our color correction system can be applied to several important applications. The first application is objective skin color processing in burn scars (Section III-G.I) and skin lesion detection (Section III-G.II). We also demonstrate how color-corrected texture maps are used to produce photo-realistic rerendering of complex surfaces for computer graphics (Section III-H).

A. Color Correction With a Single Illuminant

Fig. 7(a) is an image of three yellow (not shown in this print version) rectangular blocks on a table. We use a the PCT/MC (principal coordinate transform/median cut) algorithm to segment the surface of different colors in the scene [24]. This segmentation iteratively divides a 3-D space that is defined by the principal coordinate transform of the color information. The principal coordinate transform finds the two orthogonal directions along which the variation of the color information is largest, and then picks a third axis that completes a left-handed

coordinate frame. The general applicability and straightforward implementation of this algorithm makes it a reasonable choice for demonstration of our experimental results.

A segmentation of the original color image [Fig. 7(c)] is not able to classify all of the vertical faces of the blocks as having the same color because of the surface shading due to the varying orientations of these surfaces with respect to the illumination direction. After we applied the lighting calibration and color correction procedures to this image, these surfaces appear much more similar in color and the shading is mostly removed [see Fig. 7(b)]. The resultant segmentation [in Fig. 7(d)] is now able to correctly label all of the vertical faces as similar in color.

B. Color Correction With Multiple Illuminants

Fig. 8 shows images of southwestern native American style pottery taken with the Cyberware scanner and the corrected color images (not shown in this print version). The shading that occurred due to the uneven lighting and changing surface orientation of the pot surface has been removed, producing a color image that is suitable for photo-realistic, texture-mapped 3-D graphical rendering, as we will discuss in Section III-H.

C. Color Correction Accuracy

We would like to evaluate the performance of color correction to see how it varies with changing geometry in the scene. To test this, we corrected the color of a colorful box of compact discs placed at many different angles of incidence from the camera's viewing direction. Fig. 9 shows the original and corrected images for each orientation, which varied from 0 to 80 degrees in 10 degree steps.

TABLE I

ANALYSIS OF THE VARIANCE (ANOVA) OF (a) HUE, (b) SATURATION, AND (c) BRIGHTNESS FOR UNCORRECTED VERSUS CORRECTED COLOR PATCHES. SS IS SUMS OF SQUARES, df IS DEGREES OF FREEDOM, MS IS MEAN SQUARED, F IS THE RATIO OF TWO INDEPENDENT ESTIMATES OF THE VARIANCE (THE TWO MS VALUES), AND p IS THE PROBABILITY THAT THE VALUES OF THE VARIATES WOULD OCCUR STRICTLY BY CHANCE

| (a) | | | | |
|----------------------------|--------|----|--------|--------|
| ANOVA of hue ($p=0.867$) | | | | |
| Source | SS | df | MS | F |
| Between groups | 0.0042 | 1 | 0.0042 | 0.0287 |
| Within groups | 3.822 | 26 | 0.147 | |
| Total | 3.826 | 27 | | |

| (b) | | | | |
|-----------------------------------|----------|----|----------|----------|
| ANOVA of saturation ($p=0.997$) | | | | |
| Source | SS | df | MS | F |
| Between groups | 1.595e-7 | 1 | 1.595e-7 | 1.412e-5 |
| Within groups | 0.294 | 26 | 0.011 | |
| Total | 0.294 | 27 | | |

| (c) | | | | |
|-----------------------------------|-------|----|-------|-------|
| ANOVA of brightness ($p=0.002$) | | | | |
| Source | SS | df | MS | F |
| Between groups | 0.822 | 1 | 0.822 | 11.66 |
| Within groups | 1.832 | 26 | 0.705 | |
| Total | 2.653 | 27 | | |

There are two results from these images that are particularly interesting. First, in the original image, Fig. 9(i), one face of the box is so dark that no color is visible at all. In its corrected counterpart, Fig. 9(j), the color is made visible. Also, in Fig. 9(k), the orientation of the plane is unique in that it creates a specular reflection over almost the entire surface. Since no color correction can be applied there, Fig. 9(l) shows the result of the specular highlight prediction method, which nicely locates the highlight region.

We would like to determine whether there is any variation added to the color data by the color correction process that is statistically significant. To show this in a way that meaningfully correlates with color perception, we first convert the color to the HSI color space. In this space, we can determine whether the variation in the brightness is significant independently of the dimensions of hue and saturation, whereas in the original RGB space all three dimensions correlate with brightness. We compare patches from the images in Fig. 9, where we compute the mean hue, saturation, and brightness of 14 pairs of corresponding corrected and uncorrected regions. We use ANOVA to determine whether the difference in the variance for the hue, saturation, and brightness is significant for the corrected and uncorrected color patches. In the following ANOVA tables, SS is sums of squares, df is degrees of freedom, MS is mean squared, F is the ratio of two independent estimates of the variance (The two MS values), and p is the probability that the values of the variates would occur strictly by chance.

Table I(a)–(c) show the analysis of variance (ANOVA) for the hue, saturation, and brightness, respectively. Also, Figs. 10–12 show the distributions of hue, saturation, and brightness, respectively for uncorrected (column 1) and corrected (column




Fig. 10. Distributions of hue for uncorrected (column 1) and corrected (column 2) color patches.




Fig. 11. Distributions of saturation for uncorrected (column 1) and corrected (column 2) color patches.

2) color. There are a number of interesting analyzes we can make from this. First, the p values for hue (0.867) and saturation (0.997) support the null hypothesis, which states that these color samples are drawn from the same distribution. In other words, no statistically significant variance in hue or saturation was added as a result of the color correction process. This is especially important, since our goal is to establish an objective color measure: we do not want to add any additional “noise” to the image that would make an objective measure unreliable. Although the actual group mean hue of the uncorrected (0.439) versus corrected color (0.463) does shift due to the illumination color correction procedure, the variance in hue is not significantly different. However, the difference in brightness variation is statistically significant ($p = 0.002$). This is exactly what we expect from our procedure of reducing the variation of intensity in the image.




Fig. 12. Distributions of brightness for uncorrected (column 1) and corrected (column 2) color patches.




Fig. 13. (a) Image of four spheres with highlights. (b) Estimated positions of highlights using light source information (in white).

D. Locating Specularities

Fig. 13(a) shows an image of four colored hemispheres against a black background. We can test for specular highlights in the image using the light calibration procedure and the surface normal estimation from range data, in this case acquired with a K²T structured light scanner. The white areas in Fig. 13(b) show the areas where the incoming and outgoing light vectors are approximately equal and coplanar with the surface normal, indicating potential highlights if the surface is assumed to be specular. Alternatively, if range data is not available, then given the locations of the specular highlights in the image and an estimate of the light source direction, it is possible to estimate the surface normal at the specular locations. Thus, light source knowledge enhances the information that can be derived from specularities.

Fig. 14 shows some results from images taken with the Cyberware scanner. The images depict two brightly colored, shiny plastic dolls. When we apply color correction to these images, they take on a more cartoonish look as the shading is removed. However, the specular highlights cannot be corrected since there is little or no body reflectance information to correct. When we apply highlight detection to the images, we can see regions corresponding to the highlights. Highlights as small as one pixel in the image can be detected using these images.




Fig. 14. Results of color correction and highlight detection on Cyberware images. (a) and (d) Original color images of the shiny plastic dolls. (b) and (e) are the corresponding corrected color images. (c) and (f) Possible highlight regions in white. A supplementary color .jpg file is available at <http://ieeexplore.ieee.org>




Fig. 15. Images of (a) a flat planar surface, (b) a vertically ribbed planar surface, and (c) a bumpy, knitted planar surface.

E. Range Texture

As a first experiment for evaluating range texture, we scanned three planar patches of varying degree of coarseness. The first surface was a completely smooth block [Fig. 15(a)], the second was a cloth with a vertically-ribbed texture stretched over a block [Fig. 15(b)], and the third was a knitted sweater stretched over a block [Fig. 15(c)]. We calculated the residual image from the planar fit and computed co-occurrence matrix texture measures on four patches of size 50×50 from each of the three range images. These texture features are listed in Table II. The co-occurrence parameters were set at two pixels distance and 0° direction. For all these examples, a Bayesian linear classifier was able to distinguish the three surfaces with 100% accuracy, using six of the samples to train and six of the samples to test. Note that the variation of these feature values from each particular sub area do not vary by a large enough amount to affect classification.

Although these results are based on roughly planar surfaces, this approach is generally applicable to nonplanar surfaces and surfaces in arbitrary orientations. We have also evaluated the use of range texture in the practical application of burn scar surface characterization for color correction as detailed in [1].

TABLE II
FEATURE VALUES COMPUTED FROM FOUR DIFFERENT SUBIMAGES IN THREE
RANGE IMAGES. NOTE THAT THE FEATURE VALUES CLUSTER CLOSELY
TOGETHER IN FEATURE SPACE: A DESIRABLE QUALITY FOR INPUT TO A
CLASSIFICATION OR SEGMENTATION TASK

| Surface | Feature Name | Feature Values |
|---------|-------------------|---|
| smooth | Energy | 0.1352, 0.1624, 0.1510, 0.1770 |
| ribbed | Energy | 0.0533, 0.0403, 0.0368, 0.0272 |
| bumpy | Energy | 0.01255, 0.0110, 0.0120, 0.0139 |
| smooth | Entropy | 2.2590, 2.1345, 2.1590, 2.0150 |
| ribbed | Entropy | 3.3758, 3.6264, 3.7815, 3.9244 |
| bumpy | Entropy | 4.7648, 4.8156, 4.7098, 4.6300 |
| smooth | Local Homogeneity | 0.6316, 0.6654, 0.6669, 0.6861 |
| ribbed | Local Homogeneity | 0.5264, 0.5245, 0.5472, 0.5495 |
| bumpy | Local Homogeneity | 0.4911, 0.4433, 0.4688, 0.4678 |
| smooth | Correlation | 1.64×10^3 , 1.65×10^3 , 1.12×10^3 , 1.18×10^3 |
| ribbed | Correlation | 1.55×10^4 , 2.63×10^4 , 6.33×10^4 , 7.72×10^4 |
| bumpy | Correlation | 6.34×10^5 , 5.84×10^5 , 5.33×10^5 , 4.22×10^5 |

F. Avoiding Coarse Areas, Noise, Edges, and Highlights

To demonstrate the entire system for color correction working in concert, we examine an image from the Cyberware scanner of the author's hand and arm grasping a tall, shiny mug (Fig. 16). The long-sleeved sweater on the arm is a knitted texture that we do not want to correct, since the surface normals of the surface are not smooth. The mug is shiny and produces a specular highlight that we want to detect to alert the user and not attempt to remove shading from it. Additionally, the complex shape of the hand presents a number of edges that we wish to detect and avoid.

Fig. 16(b) shows the areas of the image (in white) where there was no 3-D information captured by the scanner. There were no objects in the background of this image that lay close enough to the scanner to reflect back the laser stripe with sufficient intensity to triangulate their position. Fig. 16(c) shows the edges (in white) of the fingers that were detected by the University of Bern range image segmenter [23], [25]. The additional white pixels in the edge image lay on the sweater, and the edge between the sweater and the hand. Fig. 16(d) shows the sweater region that was segmented and then identified as having a highly coarse texture compared to the hand and mug regions. The cup and hand regions had a lower entropy texture measure than the sweater, indicating that these surfaces are generally smoother. In other words, there is less variability in the residuals from the surface fitting on the relatively smooth hand and mug surfaces than on the bumpy sweater surface. The sweater region itself does not




Fig. 16. Example output of the complete color correction system. (a) The original color image. (b) The areas where no 3-D data was acquired, in white. (c) The edges from the range image. (d) The coarse texture region. (e) The specular highlight region. (f) Combination of (b)–(e). (g) The corrected color image. A supplementary color .jpg file is available at <http://ieeexplore.ieee.org>.

include the entire sweater surface: there is an area at the top of the sleeve that is not included in the texture mask image. This is a result of the range image segmentation that was used to identify the region. The sleeve was initially over-segmented into a larger region, the one seen in the result, and several smaller regions at the top of the sleeve. The size of the smaller regions was so small that they were ignored in the resulting segmentation. To improve the segmentation results to capture the entire sleeve would require either exhaustive tuning of the segmentation parameters for best results, or modification of the segmentation paradigm used to merge regions in order to avoid over-segmentation. Fig. 16(e) shows the specular highlight region on the cup surface that was identified using our highlight location method. There are also small regions on the top finger and the hand that are identified as having geometry to support a specular reflection. These regions are also brighter than the surrounding skin color in the image, but the presence of hair combined with the fact that skin is less specular than the plastic mug make these specularities less obvious. Fig. 16(f) shows the combination of all of the regions we have classified as areas to avoid in color correction. Finally, Fig. 16(g) shows the resulting color-corrected image (not shown in this print version). The linearization of camera response and the associated white-balancing effect was applied over the whole image, as is appropriate. Also, the correction for variation over the spatial distribution of the illumination field was applied over the entire image, which is also appropriate. Distance variation correction does not yield noticeable results in this image because all of the surfaces are very near to one another in terms of depth from the camera. We see that the hand and mug regions have the shading removed, but the




Fig. 17. Rendering of a burn scar patient (a) before color correction and (b) after color correction.




Fig. 18. (a) Color image of an arm with skin lesions. (b) Skin image patch after correction. (c) Segmentation based on image in (a). (d) Segmentation of skin lesions based on image in (b).

sweater does not, since it is identified as not being smooth. The correction of shading was also not applied at the specular highlights or the edges. There is a noticeable brightness variation on the red strip on the right of the mug that is due to inaccurate 3-D range acquisition there: the surface does not conform exactly to a cylinder there, but it slightly warped, causing the surface normal estimation to skew and over-correct for the shading.

G. Skin Color Correction

One practical application of color correction is in the objective measurement skin color to assist in medical diagnosis of burn scar or for the detection of skin lesions, including melanoma. Characterization of skin in terms of its color and texture, as captured by cameras, is confounded by lighting and

TABLE III
SKIN LESION DETECTION PERFORMANCE WITH UNCORRECTED
AND COLOR CORRECTED IMAGES

| Image Num | # of ground truth lesions | Uncorrected | | Corrected | |
|--------------|---------------------------------|--------------------|--------------------|--------------------|--------------------|
| | | False positives | False negatives | False positives | False negatives |
| 1 | 8 | 0 | 0 | 1 | 0 |
| 2 | 12 | 1 | 4 | 0 | 0 |
| 3 | 4 | 1 | 3 | 0 | 0 |
| 4 | 3 | 4 | 0 | 1 | 0 |
| 5 | 3 | 2 | 0 | 1 | 0 |
| 6 | 5 | 0 | 4 | 1 | 0 |
| 7 | 4 | 1 | 1 | 1 | 0 |
| 8 | 4 | 0 | 2 | 1 | 0 |
| 9 | 5 | 1 | 2 | 1 | 0 |
| 10 | 4 | 0 | 1 | 0 | 0 |
| 11 | 4 | 0 | 3 | 0 | 0 |
| Total | 56 | 10 | 20 | 7 | 0 |

shading artifacts. The method outlined in this paper can be used as front-end processing to arrive at objective measures.

1) *Burn Scar*: There is a need in the field of scar therapy to establish objective measurements of scar-related features such as color, texture, pliability, and volume to augment current subjective measurements. For instance, the rating of the color a patient's scar depends on the variability between human observers, the variation and color of lighting conditions surrounding the area of observation, etc. When comparing subtle distinctions in the quality of results between two different healing modalities that aim for the same results, in this case returning the skin to its original color, it is impossible to know whether the variability of the results is due primarily to the healing methods themselves, or simply due to the variations in observing color subjectively. With an objective measure, we can evaluate how well a patient is healed independently of the particular care giver who conducts the tests.

We can use a corrected color image of a burn scar patient to produce a photorealistic rerendering of the subject. Fig. 17 shows such a rendering of a burn scar patient before and after color correction. After correction, the variation in the color is much easier to see due to the brightening effect of the shading removal. Using this rerendering technique, it is possible to see how the patient would appear in light from different directions, distances, or even colors. Since the variation in illumination is removed in the corrected color image, whatever illumination is synthetically generated by the rendering model is more realistic than when we use the original color image.

2) *Skin Lesions*: Detection of skin lesions, with the ultimate goal of screening for melanoma, is another application that can benefit from color correction. The change in skin surface geometry and the constraints on the position of lights make it difficult to design automated image analysis algorithms to robustly locate small skin lesions. Here we demonstrate the effectiveness of the color correction method in enhancing skin lesion segmentation; we test the quality of segmentation before and after correction. Another goal of this experiment is to demonstrate the

applicability of the method with a different range scanner, with its own lighting arrangement. We used the Minolta 3-D Range scanner Vivid 900, which is a triangulation-based laser range scanner.

Fig. 18(a) and (b) show the uncorrected and corrected color images of an arm with skin lesions on it. Fig. 18(c) and (d) show the segmentation results on the color images in (a) and (b), respectively. The segmentation algorithm was a simple one based on local thresholding, followed by morphological operations. Since this segmentation step is the first one in the chain of further processing to characterize and classify the lesions, it is important that the missed detection rate be low, ideally zero. Notice the missed lesion in the middle portion of the arm when working with uncorrected images. This is due to the fact that arm is cylindrical in nature and the lighting source does not provide even distribution across the surface. The corrected image has a more uniform distribution of intensity, facilitating the ready detection of all the lesions.

To quantitatively evaluate the task of detection of lesions from color corrected images, we considered 11 skin images with skin lesions from three different subjects, taken at varying distances from the camera. We manually marked the lesions in each image. In total, we have 56 skin lesions. Table III shows the false positive and false negative rates for uncorrected and corrected color images. We see that for the original image we have 36% false negative rate while with the color corrected data set the rate is 0%. We also see that the false positive rates improves from 18% for the original to 13% for the corrected images.

H. Photorealistic Rerendering

In computer graphics, in order to render an object with a heterogeneously colored surface, a technique known as texture mapping is used. Texture mapping is the process of taking a two-dimensional (2-D) color image and warping it to conform to a 3-D surface. In order to produce a photorealistic rendering of an object, the texture map (color image) must have little or no illumination variation. The rendering process must be able to control the illumination that is applied to the scene. Any shading or illumination artifacts on the texture map will result in a less realistic rendering. To acquire such an image, the illumination, scene, and viewing geometry must be carefully controlled and engineered. We can alleviate this need by applying our color correction process.

Fig. 19 shows two renderings of the author's face produced from a Cyberware scan. In the first, uncorrected image, there is significant shading, particularly around the eyes and neck areas. In the second, corrected image, these shading variations have been removed. Also note that the color has become less red and more like a brown skin tone. This is because the camera response calibration process removes the tendency of the color camera to produce a redder-than-normal image.

Fig. 20 shows two renderings of a southwestern native American style pot, lit from the right. In the original scan, the texture map is lit from the left, but the color corrected texture map removes this shading, creating an image that appears properly lit in the corrected rendering.




Fig. 19. Renderings of the author's face from a Cyberware scan (a) before correction and (b) after correction. The light source is located near the camera.




Fig. 20. Renderings of a southwestern native American style pot from a Cyberware scan (a) before correction and (b) after correction. The scene is illuminated from the right. A supplementary color Mpeg animation is available for download at <http://ieeexplore.ieee.org>.

IV. CONCLUSION

We have presented here a methodology for extracting objective color from images that are made in indoor, laboratory-like environments where the number and types of illuminants may vary, a number of different cameras and 3-D ranging devices may be used, and the types of objects to be images are largely unconstrained in shape, reflectance and texture. We have demonstrated that our method is applicable to several different applications, such as imaging texture maps for photorealistic rerendering, burn scar color correction, and skin lesion detection. Images of scenes such as these that have shading and highlights are corrected to eliminate variation of illumination and model the surface reflectance parameters of an object more accurately. The code related to this work is available at <http://marathon.csee.usf.edu/jvt>

For future work, it is important for us to expand the reflectance modeling beyond the Lambertian reflectance that we

assume when removing shading. Estimation of the degree of Lambertian versus specular reflectance for each surface may yield an even more accurate measurement of intrinsic surface color. The dichromatic surface reflectance model is one such model that applies to dielectric (nonconducting) materials. Estimating of the degree of body (Lambertian) reflectance versus surface (specular) reflectance using a model such as this would be an important addition to color, especially for the application of rating burn scars.

REFERENCES

- [1] M. Powell, "Toward objective color from images," Ph.D. dissertation, Dept. Comput. Sci. Eng., Univ. South Florida, Tampa, 2000.
- [2] M. W. Powell, S. Sarkar, and D. B. Goldgof, "A simple strategy for calibrating the geometry of light sources," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 23, pp. 1022–1027, Sept. 2001.
- [3] G. Healey and D. Slater, "Computing illumination-invariant descriptors of spatially filtered color image region," *IEEE Trans Image Processing*, vol. 6, pp. 1002–1013, July 1997.

- [4] D. K. Panjwani and G. Healey, "Markov random field models for unsupervised segmentation of textured color images," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 17, pp. 939–954, Oct. 1995.
- [5] S. Nayar and R. Bolle, "Computing reflectance ratios from an image," *Pattern Recognit.*, vol. 26, pp. 1529–1542, 1993.
- [6] M. D. Fairchild, *Color Appearance Models*. Reading, MA: Addison-Wesley, 1997.
- [7] L. Maloney and B. Wandell, "Color constancy: A method for recovering surface spectral reflectance," *J. Opt. Soc. Amer. A*, vol. 3, no. 1, pp. 29–33, 1986.
- [8] D. Brainard and W. Freeman, "Bayesian color constancy," *J. Opt. Soc. Amer. A*, vol. 14, no. 7, pp. 1393–1411, 1997.
- [9] M. D'Zmura and G. Iverson, "Probabilistic color constancy," in *Geometric Representations of Perceptual Phenomena: Papers in Honor of Tarow Indow's 70th Birthday*, M. D'Zmura, D. Hoffman, G. Iverson, and K. Romney, Eds. Hillsdale, NJ: Lawrence Erlbaum, 1994.
- [10] G. Finlayson, "Color in perspective," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 18, pp. 1034–1038, Oct. 1996.
- [11] D. Forsyth, "A novel algorithm for color constancy," *Int. J. Comput. Vis.*, vol. 5, no. 1, pp. 5–36, 1990.
- [12] E. Land, "Recent advances in retinex theory and some implications for cortical computations: Color vision and the natural image," *Nat. Acad. Sci.*, vol. 80, pp. 5163–5169, 1983.
- [13] S. Shafer, "Using color to separate reflection components," *Color Res. Applicat.*, vol. 10, pp. 210–218, 1985.
- [14] S. Tominaga, "Surface reflectance estimation by the dichromatic model," *Color Res. Applicat.*, vol. 21, no. 2, pp. 104–114, 1996.
- [15] S. Tominaga and B. Wandell, "Standard surface-reflectance model and illuminant estimation," *J. Opt. Soc. Amer. A*, vol. 6, no. 4, pp. 576–584, 1996.
- [16] B. Horn, *Robot Vision*. New York: McGraw-Hill, 1986.
- [17] M. Levoy, K. Pulli, B. Curless, S. Rusinkiewicz, D. Koller, L. Pereira, M. Ginzton, S. Anderson, J. Davis, J. Ginsberg, J. Shade, and D. Fulk, "The digital michelangelo project: 3-D scanning of large statues," in *Proc. 27th Int. Conf. Computer Graphics Interactive Techniques SIGGRAPH*, 2000, pp. 131–144.
- [18] R. Jain, K. Rangachar, and B. G. Schunck, *Machine Vision*. New York: McGraw-Hill, 1995.
- [19] S. Arya, D. DeMenthon, P. Meer, and L. S. Davis, "Textural analysis of range images," *Pattern Recognit. Lett.*, vol. 12, pp. 619–626, Oct. 1991.
- [20] A. Waksman and A. Rosenfeld, "Sparse opaque three-dimensional texture, 2b: Photometry," *Pattern Recognit.*, vol. 29, no. 2, pp. 297–313, 1996.
- [21] J. Macedo, R. Manduchi, and L. Matthies, "Ladar-based discrimination of grass from obstacles for autonomous navigation," in *Experimental Robotics VII*. New York: Springer-Verlag, 2001.
- [22] M. W. Powell, S. Sarkar, and D. B. Goldgof, "Calibration of light sources," *Comput. Vis. Pattern Recognit.*, pp. 263–269, 2000.
- [23] X. Jiang, H. Bunke, and U. Meier, "High-level feature-based range image segmentation," *Image Vis. Comput.*, vol. 18, pp. 817–822, 2000.
- [24] W. V. S. Scott, E. Umbaugh, and R. H. Moss, "Automatic color segmentation of images with application to detection of variegated coloring in skin tumors," *IEEE Eng. Med. Biol. Mag.*, vol. 8, pp. 43–52, Dec. 1989.
- [25] X. Jiang and H. Bunke, "Fast segmentation of range images into planar regions by scan line grouping," *Mach. Vis. Applicat.*, vol. 7, no. 2, pp. 115–122, 1994.



Mark W. Powell received the B.S. and M.S. degrees in computer science and the Ph.D. degree in computer science and engineering all from the University of South Florida, Tampa, in 1992, 1997, and 2000, respectively.

He has been a Member of the Technical Staff in the Mobility Systems Concept Development Section at the Jet Propulsion Laboratory (JPL), Pasadena, CA, since 2001. At JPL, his area of focus is science data visualization and science planning for telerobotics.

He is a Member of the MER mission operations as

a Science Downlink Coordinator, facilitating the timely downlink and assisting with the analysis of data from the Mars Exploration Rovers.

Dr. Powell received the award for Outstanding Dissertation in the area of advanced illumination modeling; color and range image processing applied to robotics and medical imaging from the University of South Florida. He received the Imager of the Year award from *Advanced Imaging Magazine* for his work on the Science Activity Planner science visualization and activity planning software used for MER operations, and also for Maestro, the publicly available version of that same application.




Sudeep Sarkar (M'93) received the B.Tech. degree in electrical engineering from the Indian Institute of Technology, Kanpur, India, in 1988. He received the M.S. and Ph.D. degrees in electrical engineering, on a University Presidential Fellowship, from The Ohio State University, Columbus, in 1990 and 1993, respectively.

He has been with the Department of Computer Science and Engineering, University of South Florida (USF), Tampa, since 1993, where he is currently a Professor. His research interests include perceptual

organization in single images and multiple image sequences, biometrics, gait recognition, color-texture analysis, and performance evaluation of vision systems. From 2000 to 2001, he served on the editorial board of *Pattern Analysis and Applications Journal* and he is currently serving on the editorial board for *Pattern Recognition Journal*.

Prof. Sarkar was the recipient of the National Science Foundation CAREER Award in 1994, the USF Teaching Incentive Program Award for undergraduate teaching excellence in 1997, the Outstanding Undergraduate Teaching Award in 1998, and the Theodore and Venette Askounes-Ashford Distinguished Scholar Award in 2004. From 1999 to 2003, he served on the Editorial Board for the IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE. He is presently serving on the editorial boards for IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS—PART B.



Dmitry B. Goldgof (S'84–M'87–SM'93) received the Ph.D. degree in electrical engineering from the University of Illinois at Urbana-Champaign in 1989.

He is currently a Professor in the Department of Computer Science and Engineering, University of South Florida, Tampa, and a member of the H. Lee Moffitt Cancer Center. His research interests include motion analysis, computer vision, image processing and its biomedical applications, pattern recognition and bioinformatics. He has graduated ten Ph.D. and 27 M.S. students, and has published over 50 journal, over 100 conference publications, 16 books chapters, and four books. He is actively involved with a number of professional conferences. He is the North American Editor for *Image and Vision Computing Journal*.

Prof. Goldgof is an Associate Editor for IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS—PART B.

Krassimir Ivanov (M'00–A'00) received the M.S. degree in computer science from the University of South Florida, Tampa, in 2003.

He has developed and led the expansion of several server-client web based applications in his current position as Software Engineer III at Verizon Services Group, Tampa, FL. His research interests include image processing, artificial intelligence, and networking.